AFRL-IF-RS-TR-2007-58
Final Technical Report
March 2007

MULTI-SCALE BEHAVIORAL MODELING AND ANALYSIS PROMOTING A FUNDAMENTAL UNDERSTANDING OF AGENT-BASED SYSTEM DESIGN AND OPERATION

University of Texas at Austin

Sponsored by Defense Advanced Research Projects Agency
DARPA Order No. K549/01

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

NOTICE AND SIGNATURE PAGE

AFRL-IF-RS-TR-2007-58 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/ /s/

JAMES M. NAGY   JOSEPH CAMERA, Chief
Work Unit Manager   Information & Intelligence Exploitation Division
                            Information Directorate

REPORT DOCUMENTATION PAGE

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
MAR 2007	Final	Jul 00 – Sep 06

4. TITLE AND SUBTITLE

MULTI-SCALE BEHAVIORAL MODELING AND ANALYSIS PROMOTING A FUNDAMENTAL UNDERSTANDING OF AGENT-BASED SYSTEM DESIGN AND OPERATION

5a. CONTRACT NUMBER

5b. GRANT NUMBER
F30602-00-2-0588

5c. PROGRAM ELEMENT NUMBER
62301E

6. AUTHOR(S)

K. Suzanne Barber

5d. PROJECT NUMBER
TASK

5e. TASK NUMBER
00

5f. WORK UNIT NUMBER
12

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Texas at Austin
101 E 27th Street
Austin TX 78712-1500

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFRL/IFED   Defense Advanced Research Projects Agency
525 Brooks Rd   3701 North Fairfax Drive
Rome NY 13441-4505   Arlington, VA 22203-1714

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-IF-RS-TR-2007-58

12. DISTRIBUTION AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 07-084

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This research project used a combination of theory and experimentation to advance both the design methods and delivered functionality of Multi-Agent Systems in the context of military-relevant problem domains. The experiments conducted tested new theories about how systems of agents acting as distributed decision-makers should: (1) reorganize by allocating decision-making control to maximize system performance, (2) assess the trustworthiness of information by determining the level of information uncertainty and reliability of information source, and (3) coordinate by exchanging their preferences for actions under varying organizations and situations.

15. SUBJECT TERMS
Multi-Agent System (MAS), decision-making control, design methods, systems of agents, distributed decision-maker, information uncertainty, competency-based MAS designs and Sensible Agent architecture

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**	UL	25	James M. Nagy
U	U	U			**19b. TELEPHONE NUMBER** *(Include area code)*

Table of Contents

List of Figures

1.0 Summary

The research project utilized a combination of theory and experimentation to advance both the design methods and delivered functionality of Multi-Agent systems in the context of military-relevant problem domains. To date, potential users of agent technology have not had a disciplined means for designing an agent-based solution to address a given problem and assessing the value the generated solution with regard to the needs of that problem. The design and verification methods resulting from this research provided (i) a classification and metrics framework for comparing agent-based solutions/systems and their respective capabilities, (ii) tools and techniques for designing and evaluating architectures for agent-based systems, and (iii) tools and techniques for verifying that a given solution delivers intended behavior and is faithful to the architecture specifications. The research contributed advances in the theory and application of specific capabilities that enable a system of agents acting as distributed decision-makers to (i) organize the decision-making and information sharing connectivity to maximize system performance, (ii) assess the trustworthiness of information by determining the level of information uncertainty and reliability of information source and (iii) plan and coordinate efficiently by exchanging their preferences for actions under varying organizations and situations.

2.0 Introduction

This research effort aimed to transform Agent Design from a "black art" to a science. Prior Multi-Agent System (MAS) development has focused on the advancement of "Core Competencies" giving an agent the ability to: proactively plan to achieve goals or react to events [45, 54], model its environment [14], sense and act on its environment [58], communicate with other agents [53], coordinate with other agents [39, 51], and resolve conflicts [1].

MAS are systems whose BEHAVIOR is driven by those of its constituent agents (software components) which are, in turn, affected by their constituent Core Competencies and the techniques implementing those Core Competencies (Figure 1). Note: Agents may form one or more societies/organizations within the "system" where the "system" boundaries may range from the Internet to a military command-and-control system.
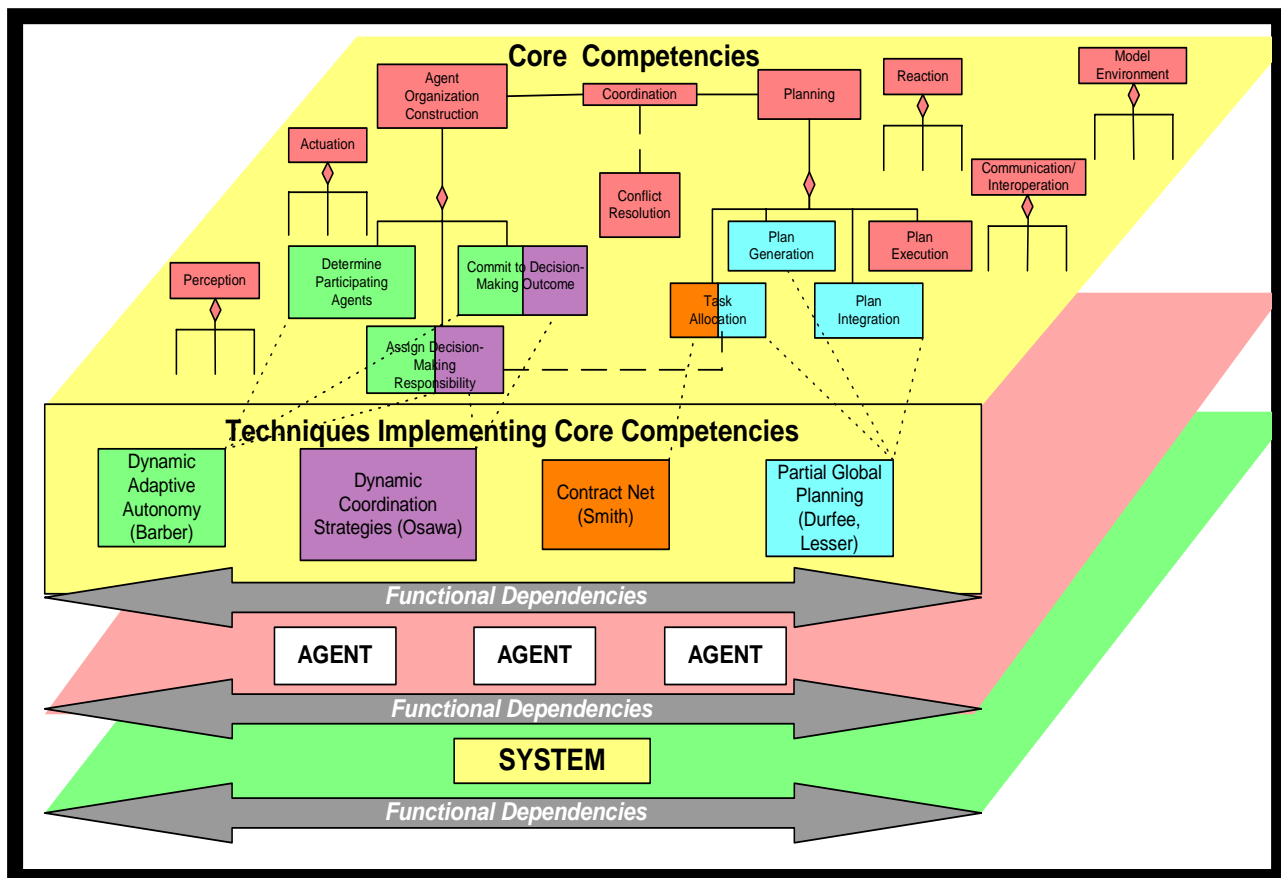

Figure 1: Core Competencies influence agent behavior and thus system behavior

The research team envisions MAS design as a two-step process:

1) Determination of the Core Competencies (behaviors) a respective agent should perform given a set of criteria mandated by the application problem domain.
2) Determination of how these selected Core Competencies will be implemented (selection of Core Competency techniques such as contract nets for coordination).

In current practice, the design process does not necessarily follow the order mentioned above. A designer may make implementation decisions first (e.g. KQML [53] for communication language, negotiation or specifically Contract Nets [61] for coordination, Partial Global Planning [40] for planning) where the decisions regarding the selection of the agent's Core Competencies are implicit. Of course, integration of the respective Core Competency techniques (implementations) is left to a skilled, well-informed agent designer and in some cases may be impossible. Many (most) MAS are developed in compliance with a specific agent design [38, 52, 57, 62] where the Core Competencies are pre-determined and technique choices are also pre-determined but may evolve in sophistication. Empirical-based, application-specific analysis has been the norm for MAS evaluation at one or multiple levels:

1) Behaviors exhibited by Core Competency techniques such as the performance of a particular coordination protocol,
2) Agent behaviors resulting from its constituent selected techniques, and
3) Overall system behavior emerging from the collective and interacting agent behaviors (Figure 1).

This typical approach to MAS design, development, and analysis severely limits the ability of MAS designers to perform bottom-up analysis; involving investigations of agent-level and system-level behavior as a function of selected Core Competency techniques, and top-down analysis; determining the effect of given agent- and system-level performance criteria on the selection of Core Competency techniques. This current practice exists for a number of reasons:

1) Designers do not have any formal models to describe MAS behaviors, behavioral interactions between these Core Competencies, and the variants of behaviors or interactions as a function of the different combinations of Core Competencies, numbers of agents, etc.
2) Designers do not have a library of possible Core Competency techniques or a formal models and analytic methods to correlate potential designs to desired behaviors.
3) Designers do not have analytical methods or tools to assist them in evaluating the complexity of both their designs, (i.e. the multitude of behaviors, interactions) and variants.

The typical ad-hoc agent design and analysis methods result in "stove-pipe" MAS solutions and research communities focusing on particular Core Competencies and techniques, as well as difficulty in migrating techniques across MAS designs or leveraging lessons learned. Promotion

of a specific Core Competency technique often ignores the question: "What is the <u>best</u> set of behaviors and corresponding implementations for this particular application?" This research worked to **change this current scenario.** We believe that the scientific formulation of multi-agent design and analysis theory could be pursued using two distinctly different methodologies:

1) Deployment of a case study research effort to evaluate a wide range of existing MAS designs, formally specify those designs then evaluate the designs according to some defined metrics (e.g. performance) to develop a repository of lessons learned – an approach to document the as-is but not learn the 1st principles of agent design and engineering to discover why behaviors exist and how new behaviors, dynamics or optimization paths can be discovered, OR
2) Specify and analyze the underlying MAS behaviors and the dependencies between those behaviors independent of decisions regarding the techniques to implement those behaviors then evaluate constraints imposed on those behaviors as a result of design decisions.

We pursued the second one. The intent of this research was to analyze and develop formal models and analysis of respective MAS behaviors (at all three levels in Figure 1) and the interaction dependencies between those behaviors to provide designers with a set of 1st principles by which they can select Core Competencies and assign techniques to realize the desired agent behaviors. This research first developed the models, analytical methods and tools giving designers the ability to construct, analyze, and discover functional/behavioral MAS architectures.

The proposed research launched from the observation: "Design decisions are intricately related to objective performance criteria and available, underlying Core Competencies as well as the techniques deployed to deliver those Core Competencies." The evaluation approach correlated design-based constraints on parameters associated with system-, agent- and Core Competency- level behaviors to analyze the resulting dynamics.

As an outcome to this funded research project, the work delivered in the following advances to the scientific community:

- Advancement #1: Theory and tools to design and verify competency-based MAS architectures

- Advancement #2: Core Competency Advances and Experimentation as implemented in the Sensible Agent architecture developed at UT-Austin

- Advancement #3: Application-related Demonstrations and Experimentations to evaluate competency-based MAS designs as well as specific core competencies.

4

3.0 Methods, Assumptions and Procedures

The proposed technical approach delivered rigor to the design and engineering of Multi-Agent Systems (MAS) by providing (1) mathematical models of MAS capabilities, functionality and operational structure, (2) formals methods to analyze and optimize MAS dynamics, and (3) an automated tool (DACAT) to leverage the mathematical models, analytic techniques and design optimization guidelines to evaluate design decisions (constraints) with regard to <u>what</u> capabilities and topologies exist in the functional/behavioral MAS architecture and <u>how</u> (implementation techniques of core competencies) the MAS system architecture is realized.

Multi-Agent System (MAS) development has focused on the advancement of "Core Competencies." We view the set CC_{Ai} of Core Competencies of agent A_i as comprised of 8-tuples involving the following basic parameters: (1) Planning proactively or deliberatively, P, (2) Reacting, R, (3) Modeling self, others and/or environment, M, (4) Sensing, S, (5) Actuation, A, (6) Coordination, Coord, (7) Communication, Comm, and (8) Conflict Resolution, CR. For course, the issue of modeling scale (level of detail) of Core Competencies is a research issue. We postulate that ***Agent Capabilities = T(Core Competencies at a given time)= T(CC_{Ai}, t).***

Therefore, at the highest level, the capabilities of agent A_0 are dictated by the Core Competencies that A_0 is endowed with. Roughly, we may say that the modeling issue is as follows: given a family of agents A_i, for i=1,…,n, each agent A_i performs, at each instant of time, t, some set of tasks related to its Core Competencies, $T_i =T(A_i, t),$ which require some set of resources that must be either provided to or owned by the agent A_i. For simplicity, we start the discussion of our research approach by considering four primary elements:

- <u>Task Space</u>: For a single agent A_i at time t, the Task Space, denoted by $T(CC_{Ai},t),$ depends on the structure of the Core Competencies CC_{Ai} assigned to the agent. For a given collection of agents, the Task Space, $T,$ is the union of the Task Spaces of the individual agents, and hence it spans the functional behaviors of all the agents in the system. When Core Competencies are decomposed into tasks and those task decomposed further into subtasks and so on, the Task Scale refers to the level of this decomposition.

- <u>Resource Spaces</u>: The Resource Space is composed of the resources provided to an agent through inputs to tasks, $IR(T_j , t)$ (e.g. sensed data or data received through communication) and the resources a respective agent owns at particular time, $OR(r, A_i, T_j, t).$ The parameter r represents the weights that are placed on the owned resource to indicate partial ownership in

the case of shared resources between agents to perform task T_j. The Resource Scales refer to the level of association of the resource space.

- <u>Agent Space</u>: The Agent space, A(*t)* is the collection of agents in the system at each instance of time, *t*. The Agent Scales within this space represent the agent organizations formed to solve goals. In other words, corresponding to goal G_i, AS(G_i,t) denotes the collection and organizational structure of agents that have teamed up to achieve this goal.

- <u>Connectivity between spaces</u>: The Connectivity is simply (but not simple) an explicit representation of the inter-dependencies between the Task, Resource and Agent spaces and scales within those spaces. Specifically, the connectivity **C** between spaces addresses the level of complexity introduced by the dependencies among tasks for the 8 different Core Competencies CC_{kAi} , k=1,…,8 within a single A_i agent, which is denoted by **C**(T*(A_i,CC_{1Ai},t),…* ,T*(A_i,CC_{8Ai},t))*. Similarly, in the case where the dependencies are between tasks and resources within a single agent this connectivity relation is denoted by **C***(T(A_i ,$CC_{1A\,i}$,t),…, T(A_i,CC_{8Ai},t), IR(T_j ,t),OR(r, A_i , T_j , t))*. Also, these relational dependencies change when the number of agents increases and/or the organizational relationships between agents change. In addition, connectivity may be random and may occur within a Core Competency, within an agent or among agents. Connectivity may be dictated by closeness of resources or tasks in their natural topology. For instance connectivity may be a function of physical distance or any other notion (resource sharing capabilities, task similarity, etc). Connectivity is the link that holds the previous spaces together and a very fundamental focus of our research.

The research effort will deliver both semantic, ontological representations of the spaces as well as mathematical representations. The exact representation of the spaces may vary and will be a modeling choice during the research effort. For example, mathematical representations of resources and tasks may be represented as sets, probability densities, or vectors of sets. These spaces exhibit some very complex dependencies. Thus, the research approach is carefully laid out to mitigate this complexity by starting with problems exhibiting reduced dependencies between spaces and building up from there. The modeling effort can be summarized in the three phases described below.

- Phase 1 focused on modeling the constituent tasks associated with each Core Competency, the resources associated with those tasks, and the task connectivity resulting from resource dependencies among tasks. A single agent system is considered in this phase, reducing significantly the complexity inherent to a multi-agent space.
- Phase 2 introduced a homogeneous multi-agent space. By homogeneous, we mean that all agents are assigned the same Core Competencies, and consequently the same Core Competency tasks (Task space) and the same set of resources to perform those tasks (Resource space). Connectivity among agents can be fully explored.
- Phase 3 introduced the concept of heterogeneous agents where agents are composed of various combinations of Core Competencies.

Electrical circuit designers are guided by a set of first principles dictated by the known underlying behaviors of current, voltage, power, etc. These behaviors have dependencies (power = current*voltage). Circuit designers currently benefit from a wide range of tools to model designs (transistor designs) such that the design reflects the desired current, voltage, and power behaviors. Additionally, these designers have tools to analyze the dynamic characteristics of their designs (e.g. ability of the transistor to deliver power given input current) as a function of design parameters or constraints on behaviors (e.g. resistance within the transistor); thus, allowing designers to specify "optimal" design specifications. These tools are possible because the behaviors, behavioral dependencies and behavioral dynamics are well known. The Designer's Agent Creation and Analysis Toolkit (DACAT) and Tracer tool are intended to provide for MAS designers what circuit designers now enjoy – understanding of the first principles of MAS behaviors and the ability to design MAS to optimize those core competency behaviors.

4.0   Results and Discussions

The results and discussion offered by this research project are defined below and organized by overarching research advances.

- Advancement #1: Theory and tools to design and verify competency-based MAS architectures

- Advancement #2: Core Competency Advances and Experimentation as implemented in the Sensible Agent architecture developed at UT-Austin

- Advancement #3: Application-related Demonstrations and Experimentations to evaluate competency-based MAS designs as well as specific core competencies.

4.1 Advancement #1: Theory and tools to design and verify competency-based MAS architectures

The research project advanced the theory and tools to design and verify competency-based MAS architectures. Tools were developed and leveraged to office MAS designer's the ability to formally specify and evaluate their Multi-Agent Systems. Specific contributions in this area are listed below:

- *Developed the Designer's Agent Creation and Analysis Toolkit (DACAT) for use by the agent-based system designer to specify an agent Reference Architecture (RA).* The RA specifies agent functionality in a domain- and technology- independent manner to foster reuse, promote separation of concerns, accommodate multiple potential agent technologies, and provide a foundation for comparing agent-based implementations. Functionality described in the agent RA is specified in terms of UT-Austin Core Competencies, which are allocated to architecture classes in the DACAT design process. The class structure is evaluated using coupling and cohesion metrics, which are good predictors of overall system qualities such as reusability and maintainability. The resulting RA is exported and used by ACET and ICET described below.

- *Developed explanation-based methods to verify consistency between the interpretation of an agent-based solution's execution and the user's comprehension of the agents' behaviors.* Accomplishments include building an Agent Explanation Ontology consisting of familiar high-level agent concepts to be used for specifying the agents' behaviors as background knowledge for the explanation-generation engine, using the ontology to model the UAV domain to enable automated explanation generation of technologies applied to UAV surveillance, developed a method to automate some of comprehension tasks that users perform to understand agent-based systems. To aid in building the background knowledge, a heuristic-based algorithm was developed to suggest possible causal relations between agent concept.

- *Developed a tool called Tracer to support explanation-based verification methods and performed preliminary analysis and verification of agent behavior in Metron's UAV simulator.* To visualize what is happening as the agent-system simulator executes, the Tracer Tool generates causal graphs representing agent behaviors, constructed from observations of agent beliefs, goals, intentions, actions, and interaction, as well as events in the environment. The causal graph has been demonstrated on two planning algorithms used in Metron's UAV simulator. The visualization aids in quick comprehension of the agent system, which brings up insightful questions to the agent designers and developers about why the agents are behaving in a certain manner or why a particular agent performed a specific action.

- *Populated a tool, the "Technology Portfolio Manager" (TPM) with specifications of DARPA TASK agent technology to aid a designer in deciding, which technologies from DARPA TASK* Agent Technology Repository to select when building an agent design, depending on (1) core agent competencies such as planning, sensing, modeling, coordinating, etc offered by the agent technology, (2) infrastructure technological requirements and (3) ability of technology to deliver specific domain requirements posed by the UAV surveillance domain.

- *Leveraged the Application Architecture Creation Toolkit (ACET) to aid designers when assessing how well selected technologies from DARPA TASK Agent Technology Repository can be reused to construct an agent-based system design.* ACET helps to evaluate the selected against operational requirements and intended architecture structure with tasks and agent classes specified in the Agent RA and evaluate the Application Architecture with respect to coupling and cohesion matrices).

- *Leveraged the Implementation Architecture Creation Toolkit (ICET) to aid designers when assessing the deployment viability of selected technologies from DARPA TASK Agent Technology Repository.* The designer can use ICET to evaluate not only the ability of agent technologies to integrate and interoperate with one another but also the probability that these agent technologies will deploy successfully on specified deployment environments (i.e. specific computational platforms, networks).

4.2 Advancement #2: Core Competency Advances and Experimentation as implemented in the Sensible Agent architecture

The research contributed advances in the theory and application of specific capabilities that enable a system of agents acting as distributed decision-makers to (i) organize the decision-making and information sharing connectivity to maximize system performance, (ii) assess the trustworthiness of information by determining the level of information uncertainty and reliability of information source and (iii) plan and coordinate efficiently by exchanging their preferences for actions under varying organizations and situations. Specific contributions are highlighted below:

- *Defined strategy components for trusting in agent social networks.* Because agents must interact with other agents whose motivations, abilities, and strategies change over time, the researchers have developed dynamic strategies by which an agent can determine which agents to interact with in order to maximize achievement of its own goals. Strategies encompass multiple decisions concerning: 1) who to trust, 2) how much to trust, 3) toward whom to behave in a trustworthy fashion, and 4) whom to take advantage of.

- *Implemented experimentation environment for observing agent trust relationships and comparing trust technologies.* The researchers played a key, guiding role in constructing the Agent Reputation and Trust (ART) Testbed for comparing trust-related technologies. This Testbed is currently in use by an international collection of research colleagues, and is being employed to compare our trust strategies against technologies developed by other researchers.

- *Advanced agent technology to promote efficient and effective coordinated information exchange within large networks of information providers.* Specifically, the agents evaluate trustworthiness, coverage, relevance, and cost of information sources and search for the most appropriate combination of information sources.

- *Continued efforts to advance agent technology designed to improve planning and coordination efficiency by exchanging preferences for actions under varying organizations and situations:*

    (1) Devised and empirically analyzed options for value propagation in Markov Decision Processes (MDPs). Bounded-depth value propagation combined with value estimation (reward combination and reward subsumption) was used to cast the MDP computation as heuristic search, yielding an eventually optimal anytime solution. MDPs calculations were adapted for situations with dynamic tasks through Goal Addition and Goal Removal algorithms. When combining these accuracy results with the observed efficiency advantages, it appears as though the efficiency gained from the approximation algorithms may well be worth any loss in accuracy, provided some errors are tolerable in the respective domain (i.e., perfect accuracy is not critical).

    (2) Enhanced previously developed techniques to improve the efficiency of action selection through the exchange of agent intentions and preferences. Researchers modified the Metron UAV simulator to include preference/commitment sharing functionality and conducted experiments comparing system performance (measured as average time between task servicing and percentage of targets visited before their lifetime expired) under various levels of situational awareness and preference sharing among UAV agents. Four coordination techniques based on passing increasing amounts of information were compared, No Coordination, Location-based Inference, Communicated Inference, and Explicit Partitioning. Experimental results show that improving situational awareness through increased coordination improves combined UAV performance.

4.3 *Advancement #3: Application-related Demonstrations and Experimentations to evaluate competency-based MAS designs as well as specific core competencies.*

Key military domain offered real world challenges to drive the research and gauge research progress. Those application domains are listed below.

- Demonstrated DACAT and core competencies developed in Advancement #2 for two military significant domains: (1) Airlift and Transport and (2) Unmanned Aerial Vehicle Surveillance and Reconnaissance.

- Engaged in significant technology transition efforts with the U.S. Navy, U.S. Army, and Texas Department of Health identifying key biosurveillance, command-and-control and maritime domain awareness applications that would highlight UT-Austin DARPA TASK research advances.

5.0   Conclusions

The DARPA TASK program and this affiliated research effort offered initial steps to advance (1) the discipline by which MAS systems are design and evaluated and the (2) core competencies which define the sophisticated behaviors of agents to self-organize and judge the agents they elect to interact with.

This research project utilized a combination of theory and experimentation to advance both the design methods and delivered functionality of Multi-Agent systems in the context of military-relevant problem domains.

Specific toolkits (DACAT, Tracer, ICET, ACET) were developed or leveraged and thereby demonstrated the ability to formally define, compare and evaluate agent architectures for multiple military applications. The verification methods used by Tracer to comprehend real-time agent behaviors proved useful for agent architectures independent of the application domain.

In terms of the advancing key competencies of multi-agent systems, experiments tested new theories about how systems of agents acting as distributed decision-makers should: (1) reorganize by allocating decision-making control to maximize system performance, (2) assess the trustworthiness of information by determining the level of information uncertainty and reliability of information source and (3) coordinate by exchanging their preferences for actions under varying organizations and situations.

6.0 Recommendations

This research team recommends defining a grand challenge problem to harvest and advance the initial seedling research explored under this grant.

7.0   References

[1]   M. R. Adler, A. B. Davis, R. Weihmayer, and R. W. Worrest, "Conflict-Resolution Strategies for Nonhierarchical Distributed Agents," in *Distributed Artificial Intelligence II*, L. Gasser and M. N. Huhns, eds. London: Pitman Publishing, 1989, pp. 139-161.

[2]   A. Arapostathis, K. Ghosh, and S. I. Marcus, "Harnack's Inequality For Cooperative Weakly Coupled Elliptic Systems," *Communications in Partial Differential Equations*, vol. 24(9-10), pp. 1555-1571, 1999.

[3]   A. Arapostathis and S. I. Marcus, "Analysis of an Identification Algorithm Arising in the Adaptive Estimation of Markov Chains," *Mathematics of Control, Signals, and Systems*, vol. 3(1), pp. 1-29, 1990.

[4]   A. Arnold, J. A. Carrillo, and I. Gamba, "Low and High Field Scaling Limits for the Vlasov- and Wigner-Poisson-Fokker-Planck Systems," Accepted to *Transport Theory and Statis. Phys.* Expected 2000.

[5]   D. G. Aronson, C. L. A., and J. L. Vázquez, "Interfaces with a corner point in one-dimensional porous medium flow," *Comm. Pure Appl. Math*, vol. 38(4), pp. 375-404, 1985.

[6]   D. G. Aronson and L. A. Caffarelli, "Optimal regularity for one-dimensional porous medium flow," *Rev. Mat. Iberoamericana 2*, vol. 2(4), pp. 357-366, 1986.

[7]   D. G. Aronson, L. A. Caffarelli, and S. Kamin, "How an initially stationary interface begins to move in porous medium flow," *SIAM J. Math. Anal.*, vol. 14(4), pp. 639-658, 1983.

[8]   I. Athanasopoulos, L. Caffarelli, and S. Salsa, "Caloric functions in Lipschitz domains and the regularity of solutions to phase transition problems," *Ann. of Math*, vol. 2(3), pp. 143, 1996.

[9]   I. Babuska, B. Andersson, P. J. Smith, and K. Levin, "Damage analysis of fiber composites," *Computer Methods in Applied Mechanics and Engineering*.

[10]   K. S. Barber, A. Goel, T. J. Graser, T. H. Liu, R. H. Macfadzean, C. E. Martin, and S. Ramaswamy, "Sensible Agents," in *Proc. 1997 IEEE International Conference on Systems, Man, and Cybernetics*, Orlando, FL, 1997, pp. .

[11]   K. S. Barber, A. Goel, D. Han, J. Kim, T. H. Liu, C. E. Martin, and R. M. McKay, "Problem-Solving Frameworks for Sensible Agents in an Electronic Market," in *Multiple Approaches to Intelligent Systems: Proceedings of the 12th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, IEA/AIE-99, Cairo, Egypt, May 31 - June 3, 1999*, I. Imam, Y. Kodratoff, A. El-Dessouki, and M. Ali, eds. Berlin: Springer, 1999, pp. 470-479.

[12]   K. S. Barber, A. Goel, D. Han, J. Kim, T. H. Liu, C. E. Martin, and R. M. McKay, "Simulation Testbed for Sensible Agent-based Systems in Dynamic and Uncertain Environments," Accepted to *TRANSACTIONS: Quarterly Journal of the Society for Computer Simulation International, Special Issue on Modeling and Simulation of Manufacturing Systems*. Expected 2000.

[13]   K. S. Barber, A. Goel, and C. E. Martin, "The Motivation for Dynamic Adaptive Autonomy in Agent-based Systems," in *Intelligent Agent Technology: Systems, Methodologies, and Tools. Proceedings of the 1st Asia-Pacific Conference on IAT, Hong Kong, December 14-17, 1999.*, J. Liu and N. Zhong, eds. Singapore: World Scientific, 1999, pp. 131-140.

[14] K. S. Barber and J. Kim, "Constructing and Dynamically Maintaining Perspective-based Agent Models in a Multi-Agent Environment," in *Proc. Third International Conference on Autonomous Agents (Agents-99)*, Seattle, WA, 1999, pp. 416-417.

[15] K. S. Barber, T. H. Liu, A. Goel, and S. Ramaswamy, "Flexible Reasoning Using Sensible Agent-based Systems: A Case Study in Job Flow Scheduling," *Production Planning and Control*, vol. 10(7), pp. 606-615, 1999.

[16] K. S. Barber, T. H. Liu, and D. C. Han, "Agent-Oriented Design," in *Multi-Agent System Engineering: Proceedings of the 9th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW'99, Valencia, Spain, June 30 - July 2, 1999*, F. J. Garijo and M. Boman, eds. Berlin: Springer, 1999, pp. 28-40.

[17] K. S. Barber, R. M. McKay, C. E. Martin, T. H. Liu, J. Kim, D. Han, and A. Goel, "Sensible Agents in Supply Chain Management: An Example Highlighting Procurement and Production Decisions," in *Proc. 19th ASME Computers and Information in Engineering Conference, Internet-Aided Design, Manufacturing, and Commerce Technical Committee*, Las Vegas, NV, 1999, pp. CIE-9078.

[18] K. S. Barber, E. White, A. Goel, D. Han, J. Kim, H. Li, T. H. Liu, C. E. Martin, and R. McKay, "Sensible Agents Problem-Solving Simulation for Manufacturing Environments," in *Proc. AAAI SIGMAN Workshop on Artificial Intelligence and Manufacturing: State of the and State of Practice*, Albuquerque, NM, 1998, pp. 1-8.

[19] M. Bardi and I. Capuzzo-Dolcetta, *Optimal Control and Viscocity Solutions of Hamilton-Jacobi-Bellman Equations* Boston: Birkhauser, 1997.

[20] L. Caffarelli and J. L. Vazquez, *Viscosity solutions for the porous medium equation* Providence, RI: Amer. Math. Soc, 1996.

[21] L. A. Caffarelli, "Allocation maps with general cost functions. Partial differential equations and applications," *Lecture Notes in Pure and Appl. Math*, vol. 177, pp. 29--35, 1996.

[22] L. A. Caffarelli, "Boundary regularity of maps with convex potentials II," *Ann. of Math*, vol. 2(144), pp. 453--496, 1996.

[23] L. A. Caffarelli, "A note on nonlinear homogenization," *Comm. Pure Appl. Math*, vol. 52(7), pp. 829--838, 1999.

[24] L. A. Caffarelli, "The regularity of monotone maps of finite compression," *Comm. Pure Appl. Math*, vol. 50(6), pp. 563-591, 1997.

[25] L. A. Caffarelli and A. Friedman, "The one-phase Stefan problem and the porous medium diffusion equation: continuity of the solution in n space dimensions," *Proc. Nat. Acad. Sci. U.S.A.*, vol. 75(5), pp. 2084, 1978.

[26] L. A. Caffarelli and A. Friedman, "Regularity of the free boundary of a gas flow in an n-dimensional porous medium," *Indiana Univ. Math. J.*, vol. 29(3), pp. 361-391, 1980.

[27] L. A. Caffarelli and C. E. Gutiérrez, "Singular integrals related to the Monge-Ampère equation Wavelet theory and harmonic analysis in applied sciences," *Appl. Numer. Harmon. Anal.*, vol. 50, 1997.

[28] L. A. Caffarelli, S. A. Kochengin, and V. I. Oliker, "On the numerical solution of the problem of reflector design with given far-field scattering data. Monge Ampère Equation: Applications to Geometry and Optimization. (Deerfield Beach, FL July 9-13, 1997).," in *Contemporary Mathematics, 226*, L. A. Caffarelli and M. Milman, eds. Providence, RI: American Mathematical Society, 1999, pp. 13-32.

[29] L. A. Caffarelli and J. L. Vázquez, "A free-boundary problem for the heat equation arising in flame propagation," *Trans. Amer. Math. Soc*, vol. 347(2), pp. 411-441, 1995.

[30] L. A. Caffarrelli and A. Friedman, "Blowup of solutions of nonlinear heat equations," *J. Math. Anal. Appl*, vol. 129(2), pp. 409-419, 1998.

[31] J. A. Carrillo, I. M. Gamba, and C. W. Shu, "Computational macroscopic approximations to the 1-D relaxation-time kinetic system for semiconductors," Submitted to *Physica D.*

[32] C. Cercignani, *The Boltzmann Equation and its Application* New York: Springer, 1988.

[33] C. Cercignani, I. M. Gamba, J. Jerome, and C. W. Shu, "Device Benchmark Comparisons via Kinetic, Hydrodynamic, and High-Field Models," Accepted to *Computer Methods in Applied Mechanics and Engineering: Special Issue on Computer Simulation of Microelectronics Problems.* Expected 1999.

[34] C. Cercignani, I. M. Gamba, J. Jerome, and C. W. Shu, "A Domain Decomposition Method for Silicon Devices," Accepted to *Transort Theory and Statis. Phys.* Expected 1998.

[35] C. Cercignani, I. M. Gamba, and C. D. Levermore, "A Drift-Collision Balance asymptotic for a Boltzmann-Poisson System in Bounded Domains," Accepted to *SIAM Applied Analysis.* Expected 2000.

[36] C. Cercignani, I.M.Gamba, J. Jerome, and C. W. Shu, "Applicability of the high field model: An analytical study via symptotic parameters defining domain decomposition," *VLSI Design*, vol. 8(1-4), pp. 275-282, 1998.

[37] C. Cercignani, R. Illner, and M. Pulvirenti, *The Mathematical Theory of Dilute Gases*, vol. 106 Springer-Verlag, 1994.

[38] D. Cockburn and N. R. Jennings, "ARCHON: A Distributed Artificial Intelligence System for Industrial Applications," in *Foundations of Distributed Artificial Intelligence*, G. M. P. O'Hare and N. R. Jennings, eds. New York: John Wiley & Sons, Inc., 1996, pp. 319-344.

[39] K. S. Decker and V. R. Lesser, "Designing a Family of Coordination Algorithms," in *Proc. First International Conference on Multi-Agent Systems*, San Francisco, CA, 1995, pp. 73-80.

[40] E. H. Durfee and V. R. Lesser, "Partial Global Planning: a Coordination Framework for Distributed Hypothesis Formation," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 21(5), pp. 1167-1183, 1991.

[41] R. Durrett, *Lecture notes on particle systems and percolation* Wadsworth & Brooks/Cole, 1988.

[42] Fernandez-Gaucherand, A. Arapostathis, and S. I. Marcus, "Analysis of an Adaptive Control Scheme for a Partially Observed Controlled Markov Chain," *IEEE Transactions on Automatic Control*, vol. AC-38, pp. 987-993, 1993.

[43] E. Fernandez-Gaucherand, A. Arapostathis, and S. I. Marcus, "On the Average Cost Optimality Equation and the Structure of Optimal Policies for Partially Observable Markov Decision Processes," *Annals of Operations Research*, vol. 29, pp. 439-470, 1991.

[44] I. M. Gamba and J. A., "Positive solutions to singular second and third order differential equations for quantum fluids," *Arch. Rat. Mech. and An.*, 1999.

[45] M. P. Georgeff and A. Lansky, "Reactive Reasoning and Planning," in *Proc. The Sixth National Conference on Artificial Intelligence (AAAI-87)*, Seattle, WA, 1987, pp. 677-82.

[46] K. Ghosh, A. Arapostathis, and S. I. Marcus, "Controlled Switching Diffusions as Hybric Processes," in *Hybrid Systems III: Verification and Control*, vol. 1066, R. Allur, T. A. Henziger, and E. D. Sontag, eds. New York: Springer-Verlag, 1996, pp. 64-75.

[47] K. Ghosh, A. Arapostathis, and S. I. Marcus, "Ergodic Control of Switching Diffusions," *SIAM J. Control and Optimization*, vol. 35(November), pp. 1952-1988, 1997.

[48] K. Ghosh, A. Arapostathis, and S. I. Marcus, "Optimal Control of Switching Diffusions with Application to Flexible Manufacturing Systems," *SIAM J. Control and Optimization*, vol. 31, pp. 1183-1204, 1993.

[49] G. Grimmett, *Percolation* Springer Verlag, 1999.

[50] T. Y. Hou, X.-H. Wu, and Z. Cai, "Convergence of a multiscale finite element method for elliptic problems with rapidly oscillating coefficients," *Math. Comp.*, vol. 68, 1999.

[51] N. R. Jennings, "Coordination Techniques for Distributed Artificial Intelligence," in *Foundations of Distributed Artificial Intelligence*, G. M. P. O'Hare and N. R. Jennings, eds. New York: John Wiley & Sons, Inc., 1996, pp. 187-210.

[52] N. R. Jennings, T. J. Norman, and P. Faratin, "ADEPT: An Agent-based Approach to Business Process Management," *ACM SIGMOD Record 27*, vol. 27(4), pp. 32-39, 1998.

[53] Y. Labrou, T. Finin, and Y. Peng, "Agent Communication Languages: The Current Landscape," *IEEE Intelligent Systems*, vol. 14(2), pp. 45-52, 1999.

[54] Y. Makoto, E. H. Durfee, T. Ishida, and K. Kuwabara, "The Distributed Constraint Satisfaction Problem: Formalization and Algorithms," *IEEE Transaction on Knowledge and Data Engineering*, vol. 10(5), pp. 673-85, 1998.

[55] R. C. Morgan and I. Babuska, "An approach for constructing families of homogenized equations for periodic media. II. Properties of the kernel," *SIAM J. Math. Anal.*, vol. 22(1), pp. 16-33, 1991.

[56] E.-I. Osawa, "A Metalevel Coordination Strategy for Reactive Cooperative Planning," in *Proc. First International Conference on Multi-Agent Systems*, San Francisco, CA, 1995, pp. 297-303.

[57] A. S. Rao and M. P. Georgeff, "BDI Agents: from Theory to Practice," in *Proc. First International Conference on Multi-Agents Systems*, San Francisco, CA, 1995, pp. 312-319.

[58] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach* Englewood Cliffs, NJ: Prentice-Hall, 1995.

[59] A. V. Sckorokhod, *Asymptotic Methods in the Theory of Stochastic Differential Equations* Providence, R.I: American Mathematical Society, 1989.

[60] P. Shi and I. Babuska, "Analysis and computation of a cyclic plasticity model by aid of Ddassl," *Comput. Mech.*, vol. 19(5), pp. 380-385, 1997.

[61] R. G. Smith, "The Contract Net Protocol: High-level Communication and Control in a Distributed Problem-Solver," *IEEE Transactions on Computers*, vol. 29(12), pp. 1104-1113, 1980.

[62] K. P. Sycara and A. S. Pannu, "The RETSINA Multiagent System: Towards Integrating Planning, Execution and Information Gathering," in *Proc. Second International Conference on Autonomous Agents*, Minneapolis/St. Paul, MN, 1998, pp. 350-351.

8.0 Acronyms

ACET   - Application Architecture Creation Toolkit
ART   - Agent Reputation and Trust
DACAT - Designer's gent Creation and Analysis Toolkit
ICET   - Implementation Architecture Creation Toolkit
KQML   - Knowledge Query and Manipulation Language
MAS   - Multi- Agent System
MDP   - Markov decision Processes
RA   - Reference Architecture
TPM   - Technology Portfolio Manager